National Center for Computational Sciences Snapshot
The Week of August 4, 2008

<u>Invisible Means of Support</u>

Astrophysicists simulate the dark matter that cradles a galaxy

A team led by astrophysicist Piero Madau of the University of California–Santa Cruz (UCSC) has given us a glimpse into the invisible world of dark matter, performing the largest computer simulation ever of dark matter evolving in a galaxy such as the Milky Way.

Madau and his collaborators—including Juerg Diemand and Marcel Zemp, both of UCSC, and Michael Kuhlen of the Institute for Advanced Study in Princeton, New Jersey—performed the simulation on Oak Ridge National Laboratory's (ORNL's) state-of-the-art Jaguar supercomputer, dividing the galaxy's envelope of dark matter into a billion parcels and showing how they would evolve over 13 billion years. The collaborators review the simulation and their findings in the August 7 issue of the journal *Nature*. Their article is entitled "Clumps and Streams in the Local Dark Matter Distribution."

Scientists believe visible matter provides far too little gravitational force to keep stars and galaxies tethered to the orbits we observe. In fact, researchers have concluded that what we see makes up less than a fifth of the matter in the universe. The rest, known as dark matter, has no interaction with regular matter except through the force of gravity. Nevertheless, there is so much dark matter in the universe that its gravitational force controls the lives of stars and galaxies.

The simulation performed by Madau and his teammates used about 1 million processor hours, following a galaxy worth of dark matter through nearly the entire history of the universe. It was a staggering job, tracking 9,000 trillion trillion trillion tons of invisible stuff spread across 176 trillion trillion trillion square miles as it evolved over 13 billion years. Each of the billion parcels of dark matter was 4,000 times as massive as the sun.

Dark matter is not evenly spread out, although researchers believe it was nearly homogeneously distributed right after the Big Bang. Over time, however, it became bunched and clumped as gravity pulled it together, first into clumps more or less the mass of Earth. These were pulled together into larger clumps, which were pulled together into still larger clumps, and so on until they combined to form halos of dark matter massive enough to host galaxies.

One key question answered by Madau's team is whether the smaller clumps would remain identifiable or would smooth out within the larger galactic halos. In fact, the team showed that dark matter is clumped very unevenly throughout the galaxy

"You expect a hierarchy of structure in cold dark matter," Madau explained. "What you don't know is what sort of structure will survive the assembly because as these subclumps come together, they are subject to tidal forces, and they can be stripped and destroyed. So their existence in the field had been predicted. The issue was whether they would survive as assembled together to bigger and bigger structures."

"What we find," he continued, "is the survival fraction is quite high."

Because they did not have the resolution to resolve any unevenness, smaller simulations showed the dark matter smoothing out, especially in the galaxy's dense inner reaches. Madau's billion-cell simulation (five times larger than any other of its type), however, provided enough resolution to verify that subclumps and sub-subclumps do indeed survive, even in the galaxy's inner regions.

Madau's team will be able to verify its simulation results using the National Aeronautics and Space Administration's Gamma-Ray Large Area Space Telescope (GLAST). Launched on June 11, GLAST will scan the heavens to study some of the universe's most extreme and puzzling phenomena: gamma-ray bursts, neutron stars, supernovas, and dark matter, just to name a few. While dark matter particles cannot themselves be detected by GLAST (direct detection of dark matter is being pursued by large underground detectors), researchers believe that dark matter particles and antiparticles may be annihilated when they bump into each other, producing gamma rays that can be observed from space. The clumps of dark matter predicted by Madau's team should bring more particles together and thereby produce an increased level of gamma rays. In fact, Madau and his colleagues have produced a second paper, "The Dark Matter Annihilation Signal from Galactic Substructure: Predictions for GLAST," which is scheduled for publication September 10 in *Astrophysical Journal*.

A second verification comes from an effect known as gravitational lensing, in which the gravity exerted by a galaxy along the line of sight bends the light traveling from faraway quasars in the background. If the dark matter halos of galaxies are as clumpy as this simulation suggests, the light from a distant quasar should be broken up, like a light shining through frosted glass.

"We already have some data there," Madau noted, "which seems to imply that the inner regions of galaxies are rather clumpy. Instead of a smooth lens, there is substructure that appears to be affecting the lensing process. Our simulation seems to produce the right amount of lumpiness."

While this simulation may help explain why the universe is as it is, it leaves other questions unanswered. For example, Madau noted that while dark matter in the galaxy is very clumpy, that same feature is not found in visible matter. In fact, the Milky Way is orbited by far fewer dwarf galaxies than might be predicted by the simulation.

"We see as many as 20 dwarf galaxy satellites of the Milky Way that are luminous," he explained, "and we estimate that there may be a hundred or so. But we predict thousands of relatively massive clumps. So one issue is, why are most of those clumps dark? It seems that we're not very efficient at forming stars."

<u>Upgraded Jaguar's Transition to Operations Features Six Pioneering Applications</u>

Software codes exploit majority of Jaguar's muscle

Six select software applications have been running pioneering "science-at-scale" simulations on high-performance computers at ORNL. The simulations, carried out at the National Center for Computational Sciences (NCCS), employ most or all of the processing cores of the center's flagship system, a Cray XT4 supercomputer called Jaguar that was upgraded in May to perform 263 trillion calculations per second, or teraflops.

The upgrade increased Jaguar's ability to further contribute to the Department of Energy's (DOE's) Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, which facilitates some of the world's most challenging computer simulations by issuing a small number of large awards, each consisting of millions of processing hours, to premier researchers representing diverse scientific fields and institutions. In 2008 INCITE gave 30 projects more than 145 million processing hours on ORNL supercomputers.

Running computationally demanding software applications after a major machine upgrade is part of a transition-to-operations activity, dubbed T2O, that allocates up to 4.5 million hours to each application that can concurrently use the majority of Jaguar's 31,000 processing cores. As a leadership computing facility for open scientific research, the NCCS is responsible for selecting, procuring, deploying, and operating next-generation high-performance computing systems for DOE.

The center's staff collaborates intensely with industry, academia, and government to develop leadership systems for scientific simulations in fields including climate, nuclear fusion, astrophysics, nanoscience, chemistry, biology, combustion, accelerator physics, and engineering. The capacity-class simulations run at the NCCS advance America along the path to petascale computing, capable of a quadrillion calculations per second, by testing high-performance computing systems and software applications that support innovative, high-impact science.

When a commissioned NCCS system has passed a formal acceptance test, it immediately enters the T2O phase, during which time its performance is monitored and assessed. One of the first researchers to put Jaguar to the test this spring was Jacqueline Chen of Sandia National Laboratories, who ran an application called S3D that simulated key underlying processes in combustion. Her simulations of ethylene, a hydrocarbon fuel, required 4.5 million hours running on 30,000 processors and generated more than 50 terabytes of data—more than five times as much data as contained in the printed contents of the U.S. Library of Congress.

Other researchers with pioneering runs that put Jaguar through its paces were Anthony Mezzacappa, ORNL (core-collapse supernovas); Bill Tang, Princeton Plasma Physics Laboratory (fusion plasmas); Thomas Schulthess, ORNL (Mott insulators, cuprate superconductors, and nanoscale systems); Robert Harrison, ORNL (chemical catalysts); and Synte Peacock, University of Chicago (turbulent transport in the global ocean). The software applications they ran were, respectively, Chimera, GTC, DCA++, MADNESS, and POP.

Journal Cover Highlights Jaguar Simulations
ORNL-led team featured in leading plasma publication

Fusion simulations performed on ORNL's Cray XT4 Jaguar supercomputer are featured in the cover article of July's edition of the journal *Physics of Plasmas*, published by the American Institute of Physics.

A team led by ORNL physicist Fred Jaeger used its AORSA code to demonstrate that radio waves will be effective in heating the multinational ITER fusion reactor. The team's article is entitled "Simulation of high-power electromagnetic wave heating in the ITER burning plasma," and the magazine's cover features an image created by NCCS visualization specialist Sean Ahern.

The ITER reactor will use antennas to launch radio waves carrying 20 megawatts of power into the fusion plasma, an ionized gas containing deuterium and tritium. The waves will both heat the plasma—which must reach a temperature about ten times hotter than the center of the sun—and create a current that controls it. The team's simulations will help the reactor's designers configure the antennas to make the most of that power in both these areas.

The team is part of a Scientific Discovery through Advanced Computing (SciDAC) project known as the SciDAC Center for Simulation of Wave-Plasma Interactions. Its AORSA code has been especially effective at making use of Jaguar's enormous computing power, reaching 154 trillion calculations a second.

NCCS Hosts IBM Blue Gene/P Workshop
Blue Gene architect on site

The NCCS recently hosted an educational seminar on the IBM Blue Gene supercomputing platform.

Besides presenting information on the Blue Gene architecture and instructing researchers on ways to most effectively use the platform, the seminar allowed hands-on access to the NCCS's IBM Blue Gene/P. Dubbed "Eugene," the system features 8,192 compute cores and a peak performance of more 27 teraflops.

Researchers were given the opportunity to port their codes to the new architecture and then run them, which they did with success. "A number of researchers expressed interest in pursuing an official allocation on Blue Gene to utilize the machine in the future," said NCCS staff member Don Frederick.

Nearly 30 researchers attended from ORNL, the NCCS's home campus, and core participating universities, including Duke University, the University of Virginia, and Vanderbilt University.

The designer of the Blue Gene/P, Rajiv Bendale, also visited the workshop and shared his expertise. The Blue Gene/P represents the latest advancement of the Blue Gene supercomputing platform, most notable for its scalability (the hardware is optimized to scale to thousands of cores) and energy efficiency.

The conference took place from July 29–31 at the Pollard Technology Conference Center in Oak Ridge, Tennessee.

<u>Students Build World's Fastest Supercomputer, Sort Of</u>
ARC program paves way for future scientist

Engaging young people in math and science is critical to America's technological and economic future. This is exactly the goal of the Appalachian Research Council's (ARC's) Math, Science, and Technology Institute, held at ORNL from July 12–25.

The event brings students and teachers from all over Appalachia to Oak Ridge, Tennessee, for an up-close look at scientific research. This year the students and teachers were separated into 12 teams, and each team was given a specific research problem to investigate. Laboratory staff mentored the teams throughout the research process, after which each team made a presentation detailing its results.

Five students worked closely with the NCCS's Bobby Whitten and Mitchell Griffith in a research project titled "Build a Supercomputer!" The project involved working as a team and learning basic networking, UNIX, programming, parallel programming, and MPI. The team was asked to build a supercomputer using five Macintosh Minis and find out in which year the resulting system would have been the fastest in the world.

The team created a local area network, installed software, wrote a "Hello World" parallel program, ran the program, and benchmarked the assembled system. Theoretically, the peak performance of the combined Minis was 13 gigaflops, but because of the continual failure of one of the Minis during benchmarking, the resulting peak performance was only 5.8 gigaflops. Therefore, the team concluded, the system would have been the fastest machine in the world in 1990. Overall, said NCCS staff member and team mentor Bobby Whitten, the project was a success. "The team completed all objectives and made a very professional presentation," he said.

"I learned so much in such an incredibly short amount of time," said student and ARC participant Rebecca Harrison of Richmondville, New York. "It is memories like this opportunity that will last a lifetime and help so much in the future. I am actually seriously interested in coming back to Oak Ridge, whether as an intern or for a career. I just absolutely loved it. It's just giving me another goal and something else to look forward to."